

A Note on the Fairness of Additive Increase and Multiplicative Decrease

Paul Hurley *, Jean-Yves Le Boudec, Patrick Thiran
 Institute for Computer Communication and Applications (ICA)
 Ecole Polytechnique Fédérale de Lausanne (EPFL)

Some recent papers [5,9] have shown that congestion control based on additive increase and multiplicative decrease tends to share bandwidth according to proportional fairness. Proportional fairness is a form of fairness which distributes bandwidth with a bias in favour of flows using a smaller number of hops; this is in contrast with max-min fairness, which gives absolute priority to small flows. We revisit these results by using the modelling framework based on the ordinary differential equation method in [7] and [6]. We find that for the case of small increments and constant round trip times, and in the regime of rare negative feedback, the proportional fairness result can only very approximately reflect the real rate allocation when we assume that the feedback received by sources is independent of their sending rates. In the case where sources receive feedback proportionally to their sending rates, and still for sources with identical round trip times, this is no longer true and the fairness provided is different. We show, by simulation on some examples, that even for larger increments, the average rate convergence is in agreement with our results. Finally, we establish that in the event of rate proportional feedback, our results maintain consistency with the well-known derivations relating TCP throughput as a function of loss ratio. However, this does not hold for the rate independent case, which we consider further validation of the assumption of rate dependent feedback.

1. Introduction

In this article, we revisit the topic of the distribution of rates as determined by adherence to the additive increase/multiplicative decrease algorithm. This algorithm [13] was originally believed to exhibit max-min fairness, an allocation favouring smaller rates. This is the allocation reached such that any further increase in the rate of one source results in the decrease of some smaller rate. Results in [5,9] showed that for equal round-trip times TCP appeared to provide proportional fairness, a form of fairness which distributes bandwidth with a bias in favour of flows using a smaller number of hops.

We argue that TCP connections of equal round-trip times do not converge to long term rates in agreement with proportional fairness. Rather, we show that in the event of rare negative feedback and equal round trip times, TCP distributes rates more closely in accordance with the fairness distribution algorithm derived here, F_A -fairness.

*Contact author: paul.hurley@epfl.ch, Ph. +41-21-693-6626, Fax +41-21-693-6610

Even in the event where we have rate independent feedback we show a result which closer reflects the convergence than proportional fairness. To this aim, we use the tool of the method of ordinary differential equation to examine the development of long term rates for different sources. This establishes, in the event of rare negative feedback, convergence to F_A -fairness, as the multiplicative decrease and linear increase factors approach zero. We subsequently show, by simulation, that for large factors such as those specified by TCP, the average rate for each source converges around the value determined by F_A -fairness.

We demonstrate the behaviour of an F_A -fairness distribution in the context of the well-known example; the parking lot scenario. Finally, we establish that in the event of rate proportional feedback, our results maintain consistency with the well-known derivations relating TCP throughput as a function of loss ratio. However, this does not hold for the rate independent case, which we consider further validation of the assumption of rate dependent feedback.

2. Model

We consider a simplified network model, as follows. Traffic sources, labelled $1, \dots, i, \dots, I$, send data to one destination. The network is viewed as a collection of links labelled $1, \dots, l, \dots, L$, where the only resource to be consumed is link bandwidth. Every traffic source uses a fixed route. We call x_i the sending rate for source i and assume that the amount of traffic from source i carried on link l is $A_{l,i}x_i$. The latter assumption amounts to assuming that losses are negligible. If source i sends traffic to one or several destinations over one single route, then $A_{l,i} = 0$ or 1 for all l , and those links l for which $A_{l,i} = 1$ constitute the route followed by the data. The general case where $A_{l,i}$ may have values between 0 and 1 allows traffic splitting over parallel paths.

We assume that the rates of all sources are controlled by a mechanism of additive increase and multiplicative decrease as is encountered with TCP or ATM ABR.

Modelling this mechanism is very complex because it contains both a random feedback (under the form of packet loss) and a random delay (the round trip time, including time for destinations to give feedback). In this paper we consider that all round trip times are constant and all equal. In a further paper we will consider constant round trip times that are not equal for all sources. We model the system as follows.

We consider a number of time cycles or duration τ , where τ is the common round trip for all sources. During time cycle number t , the source sending rate for source i is assumed to be constant, and is noted $x_i(t)$. At the end of time cycle number t , source i receives a random, binary feedback $E_i(t)$, which is used to compute a new value of the sending rate. The binary feedbacks $E_i(t)$ for all i are independent Bernoulli random variables, conditionally to the state of the system $\vec{x}(t) = (x_1(t), \dots, x_i(t), \dots, x_I(t))$. The sequence $\vec{x}(t)_t$ is thus a markov chain. The feedback models packet losses in the Internet, or the congestion experienced bit in DecNet, Frame Relay or ATM. In this paper, we assume the regime of rare negative feedback, and thus $E_i(t)$ takes values in the set $\{0, 1\}$.

Sources react to feedback by adjusting their rate, using an additive increase when $E_i(t) = 0$ and a multiplicative decrease when $E_i(t) = 1$. This gives the following equation.

$$x_i(t+1) = x_i(t) + r_0(1 - E_i(t)) - E_i(t)(\eta x_i(t)) \quad (1)$$

or equivalently

$$x_i(t+1) = x_i(t) + r_0 - E_i(t)(r_0 + \eta x_i(t)). \quad (2)$$

In the equation, r_0 is the rate additive increment and η the multiplicative decrease factor. For TCP, ignoring the effect of exponential increase during slow start, and assuming that all packets have the same size, we have $r_0 = 1/\tau$ (in packets per second) and $\eta = 0.5$.

As discussed later, we derive a behaviour in an ideal case where, unlike with the real TCP implementations, r_0 and η are small. Afterwards we present simulation results which show that a TCP-like connection's average rate converges in agreement with our results.

We also assume that all packets have the same, fixed size, as with ATM. The amount of negative feedback received during one time cycle of duration τ is equal in average to $\mathbb{E}(E_i(t)|\vec{x}(t))$, which is the expectation of $E_i(t)$ conditionally to $\vec{x}(t)$.

We consider two possible cases for the distribution of feedback.

Case A: rate proportional feedback

The expectation of $E_i(t)$ conditionally to $\vec{x}(t)_t$ is given by

$$\mathbb{E}(E_i(t)|\vec{x}(t)) = \tau \sum_{l=1}^L g_l(f_l(\vec{x}(t))) A_{l,i} x_i(t) \quad (3)$$

with $f_l(\vec{x}(t)) = \sum_{j=1}^L A_{l,j} x_j(t)$. In the formula, $f_l(t)$ represents the total amount of traffic flow on link l , while $A_{l,i}$ is the fraction of traffic from source i which uses link l . We interpret Equation (3) by assuming that $g_l(f)$ is the probability that a packet is marked with a feedback equal to 1 (namely, a negative feedback) by link l , given that the traffic load on link l is expressed by the real number f ; in the regime of rare negative feedback, we assume that we can neglect the occurrence of one packet marked with a negative feedback on several links within one time cycle. Then Equation (3) simply gives the expectation of the number of marked packets received during one time cycle by source i .

We surmise that this models accurately the case where all flows receive the same loss rate independent of packet level statistics. This is believed to be achieved by using active queue management such as RED [3].

Case B: rate independent feedback

In this hypothetical case, the expectation of the amount of feedback received per cycle would have the form

$$\mathbb{E}(E_i(t)|\vec{x}(t)) = C \sum_{l=1}^L g_l(f_l(\vec{x}(t))) A_{l,i} \quad (4)$$

In the formula, C is a constant, and the rest is as for case A. We do not think that this case is a realistic model for congestion control under the assumption of rare negative feedback, and examine it partly because it implicitly underlies the findings in [5,9,4].

3. The method of the ordinary differential equation

With our system model, $\vec{x}(t)$ is a Markov chain and the transition probabilities can be entirely defined using Equations (2) and (3) for case A, or 2 and 4 for case B. We use here

an alternative tool, which gives some insight about the convergence of the system. The tool is the method of the Ordinary Differential Equation (ODE), which was developed by Ljung [7] and Kushner and Clark [6]. The method applies to stochastic iterative algorithms of the form

$$\vec{x}(t+1) = \vec{x}(t) + \gamma \vec{H}(\vec{\xi}(t), \vec{x}(t)) \quad (5)$$

where $\vec{\xi}(t)$ is a sequence of random inputs and $\gamma > 0$ a small gain parameter, to which we associate the ordinary differential equation

$$\frac{d\vec{x}(s)}{ds} = h(\vec{x}(s)) \quad (6)$$

$$\text{where } h(\vec{x}) = \mathbb{E}\{\vec{H}(\vec{\xi}, \vec{x}(t)) | \vec{x}(t)\}. \quad (7)$$

The result of the method is that the stochastic system in Equation (5) converges, in some sense, towards an attractor of the ordinary differential equation (O.D.E.) in Equation (6). An attractor \vec{x}^* of the ordinary differential equation is defined by the fact that the solutions $\vec{x}(t)$ of Equation (6) satisfy $\lim_{t \rightarrow +\infty} \vec{x}(t) = \vec{x}^*$ for appropriate initial conditions. We are interested here in the case where the attractor is an equilibrium point.

Here $\vec{\xi} = \vec{E} = (E_1, E_2, \dots, E_I)$. Since r_0 and η are small, we can write $r_0 = k_r \gamma$ and $\eta = k_\eta \gamma$ where k_r and k_η are two positive constants. Then $\vec{H} = (H_1, \dots, H_I)$ with $H_i(\vec{E}, \vec{x}) = k_r - E_i(k_r + k_\eta x_i)$. The components of the mean vector field $\vec{h}(\vec{x})$ are therefore,

$$h_i(\vec{x}) = k_r - \tau x_i(k_r + k_\eta x_i) \sum_{l=1}^L g_l(f_l(\vec{x})) A_{l,i}$$

in Case A and by a similar expression for Case B. As the random feedback $\vec{E}(t)$ are independent variables depending only on the latest value of $\vec{x}(t)$, and as the mean vector field satisfies the requirements of Theorem 3 of Chapter 2 from [1], we can apply this theorem, which we rephrase as follows:

Theorem 3.1 *If the ordinary differential equation (6) is globally stable, with a unique stable equilibrium \vec{x}^* , then for $\gamma > 0$ sufficiently small, for all $\varepsilon > 0$, there exists a constant $C(\gamma)$ tending towards zero as γ tends to zero, such that*

$$\limsup_{t \rightarrow \infty} \mathbb{P}\{\|\vec{x}(t) - \vec{x}^*\| > \varepsilon\} \leq C(\gamma). \quad (8)$$

Note that multiplying the right-hand side of Equation (6) by $\gamma > 0$ does not modify the convergence properties of the O.D.E. (it only amounts to a change of time scale). For simplicity of notation, we therefore study the equivalent O.D.E. $\frac{d\vec{x}(s)}{ds} = \gamma h(\vec{x}(s))$.

4. Application to the analysis of cases A and B

We apply the method of the ordinary differential equation to find some properties of our system. First we need to study the ODE for both cases.

Case A (rate proportional feedback)

Combining Equations (6), (7) with (2) and (3), we obtain:

$$\frac{dx_i}{ds} = r_0 - \tau x_i (r_0 + \eta x_i) \sum_{l=1}^L g_l(f_l) A_{l,i} \quad (9)$$

where $f_l = \sum_{j=1}^I A_{l,j} x_j$. In order to study the attractors of this ODE, we identify a Lyapunov for it [11]. To that end, we follow [5] and [4] and note that

$$\sum_{l=1}^L g_l(f_l) A_{l,i} = \frac{\partial}{\partial x_i} \sum_{l=1}^L G_l(f_l) = \frac{\partial G(\vec{x})}{\partial x_i}$$

where G_l is a primitive of g_l defined for example by

$$G_l(f) = \int_0^f g_l(u) du \quad \text{and} \quad G(\vec{x}) = \sum_{l=1}^L G_l(f_l).$$

We can then rewrite Equation (9) as

$$\frac{dx_i}{ds} = x_i (r_0 + \eta x_i) \left\{ \frac{r_0}{x_i (r_0 + \eta x_i)} - \tau \frac{\partial G(\vec{x})}{\partial x_i} \right\} \quad (10)$$

Consider now the function J_A defined by

$$J_A(\vec{x}) = \sum_{i=1}^I \phi(x_i) - \tau G(\vec{x}) \quad (11)$$

with $\phi(x_i) = \int_0^{x_i} \frac{r_0 du}{u(r_0 + \eta u)} = \log \frac{x_i}{r_0 + \eta x_i}$. Then we can rewrite Equation (10) as

$$\frac{dx_i}{ds} = x_i (r_0 + \eta x_i) \frac{\partial J_A(\vec{x})}{\partial x_i}. \quad (12)$$

Now it is easy to see that J_A is strictly concave and therefore has a unique maximum over any bounded region. It follows from this and from Equation (12) that J_A is a Lyapunov for the ODE in (9), and thus, the ODE in (9) has a unique attractor, which is the point where the maximum of J_A is reached.

Combined with Theorem 3.1, this shows that, for case A, the rates $x_i(t)$ converge at equilibrium towards a set of values that maximise $J_A(\vec{x})$, with J_A defined by

$$J_A(\vec{x}) = \sum_{i=1}^I \log \frac{x_i}{r_0 + \eta x_i} - \tau G(\vec{x}).$$

Case B (rate independent feedback)

The analysis follows the same line. The ODE is now

$$\frac{dx_i}{ds} = r_0 - C(r_0 + \eta x_i) \sum_{l=1}^L g_l(f_l) A_{l,i} \quad (13)$$

from where we derive that, for case B, the rates $x_i(t)$ converge at equilibrium towards a set of value that maximises $J_B(\vec{x})$, with J_B defined by

$$J_B(\vec{x}) = \frac{r_0}{\eta} \sum_{i=1}^I \log(r_0 + \eta x_i) - CG(\vec{x}).$$

Interpretation and Comparison with previous results

In order to interpret the previous results, we follow [5] and assume that, calling c_l the capacity of link l , the function g_l can be assumed to be arbitrarily close to δ_{c_l} , in some sense, where $\delta_c(f) = 0$ if $f < c$ and $\delta_c(f) = 1$ if $f \geq c$. Thus, at the limit, the method in [5] finds that, for case A, the rates are distributed so as to maximise

$$F_A(\vec{x}) = \sum_{i=1}^I \log \frac{x_i}{r_0 + \eta x_i}, \quad (14)$$

subject to the constraints $\sum_{j=1}^I A_{l,j} x_j \leq c_l$ for all l . For case B, the rates tend to maximise

$$F_B(\vec{x}) = \sum_{i=1}^I \log(r_0 + \eta x_i), \quad (15)$$

subject to the constraints $\sum_{j=1}^I A_{l,j} x_j \leq c_l$ for all l .

Now we compare these results with the results recalled in the introduction. Both [5] and [9] find that, under the limiting case mentioned where g_l tends to δ_{c_l} , the rates x_i are distributed according to proportional fairness. This is equivalent to stating that the rates x_i tend to maximise $F_0(\vec{x}) = \sum_{i=1}^I \log x_i$, subject to the constraints $\sum_{j=1}^I A_{l,j} x_j \leq c_l$ for all l . If we compare our results, we find two differences. Firstly, in [5] and [9], the model implicitly assumes case B, whereas we contend that case A is more realistic, in the regime of rare negative feedback.

Secondly, even for case B, our results do not exactly coincide. Indeed, in [5] and [9], the system is directly modelled with a differential equation, without using the intermediate stochastic modelling as we do in Section 3. The differential equation in [5] and [9] is

$$\frac{dx_i}{ds} = C \left(r_0 - \eta x_i \sum_{l=1}^L g_l(f_l) A_{l,i} \right)$$



which differs from Equation (13) by a missing term r_0 in the second part, and the constant C being outside. It is our interpretation that our modelling method using the stochastic system more accurately reflects the real behaviour of the additive increase, multiplicative decrease algorithm, at least for the cases where our assumptions hold.

If we compare case B versus proportional fairness, we find that, since r_0 is assumed to be small, the difference between F_B and F_0 is small, and thus, if feedback is distributed independent of the sending rate, then rates tend to be roughly distributed according to proportional fairness. In some sense, this confirms the results in [5] and [9]. However, on the example of the next section, we find that case B tends to give less to sources that use several bottleneck links.

The situation is very different for case A, which we claim is more realistic. Here, the weight given to x_i tends to $-\log \eta$ as x_i tends to $+\infty$. Thus, the distribution of rates will tend to favour small rates more than proportional fairness would. In the next section we find an example that is indeed between proportional and max-min fairness.

5. Examples of F_A and F_B Fairness

We define F_A -fairness and F_B -fairness as the distribution given by maximising F_A and F_B respectively as shown in Equations (14) and (15). In this section we show, for the

Figure 1. Parking lot Scenario with I linksFigure 2. Numerical illustration of Section 5.2: $\frac{x_0}{c}$ as a function of c .

example of the parking lot scenario, that F_A -fairness allocates more to sources that would receive a small rate allocation from proportional fairness, and less to these sources than max-min fairness. In event of very small capacity, it approximates proportional fairness. For large c , F_A -fairness varies between max-min and proportional fairness.

We also show that F_B -fairness always allocates less than proportional fairness would to sources that would get small rates from proportional fairness.

5.1. Parking Lot Scenario

The (in)famous parking lot scenario is shown in Figure 1. It consists of I links each with capacity c . Sources of type 0 traverse the entire I links, while sources of type $i \geq 1$ only traverse the i th link. The number of sources of each type is given by $\vec{n} = (n_0, n_1, \dots)$.

The distribution for max-min fairness and proportional fairness in the parking lot scenario is $x_0 = \frac{c}{n_0 + \max_{i=1, \dots, I} n_i}$ and $x_0 = \frac{c}{\sum_{i=0}^I n_i}$ respectively [9].

5.2. Analysis of F_A -fairness

Here we analyse, in the context of the parking lot scenario, the nature of rate distributions given by F_A . The fraction of capacity distributed by F_A -fairness is not independent of the capacity, unlike the proportional and max-min fairness cases.

Since $n_0 x_0 + n_i x_i = c$, F_A can be expressed in terms of x_0 ,

$$F_A(x_0) = n_0 \log \left(\frac{x_0}{r_0 + \eta x_0} \right) + \sum_{i=1}^I n_i \log \left(\frac{c - n_0 x_0}{r_0 n_i + \eta(c - n_0 x_0)} \right). \quad (16)$$

Note that $F_A(x_0)$ goes to $-\infty$ as x_0 goes to 0 and $\frac{c}{n_0}$. This guarantees that at least one maximum in the valid range, $x_0 \in (0, \frac{c}{n_0})$. We can thus determine the distribution of \vec{x} , by solving $F'_A(x_0) = 0$. For general \vec{n} maximising this directly soon becomes messy, as it involves solving a polynomial of order up to $2I$. So we focus on the case when $\vec{n} = (v, w, w, \dots)$, for which the follow result may be derived:

Lemma 5.1 *The F_A -fairness distribution for the parking lot scenario where $\vec{n} = (v, w, w, \dots)$ is given by*

$$x_0 = \frac{v(2c\eta + r_0 w) + I w^2 r_0 - \sqrt{(v(2c\eta + r_0 w) + I w^2 r_0)^2 - 4(v^2 - I w^2)c\eta(\eta c + r_0 w)}}{2\eta(v^2 - I w^2)} \quad (17)$$

when $v^2 - Iw^2 \neq 0$, and

$$x_0 = \frac{c(\eta c + r_0 w)}{Iw^2 r_0 + v(2c\eta + r_0 w)} \quad (18)$$

when $v^2 - Iw^2 = 0$. x_i is then given by $x_i = \frac{c-vx_0}{w}$, $i = 1 \dots I$.

When $\vec{n} = (v, w, w, \dots)$, the distribution for max-min fairness and proportional fairness is given by $x_0 = \frac{c}{v+w}$ and $x_0 = \frac{c}{v+Iw}$ respectively.

To examine how F_A -fairness distribution varies with c we examine the fraction of capacity source 0 receives, $\frac{x_0}{c}$, as capacity increases.

For F_A -fairness x_0/c is increasing in c and we determine from Equation (17) that,

$$\lim_{c \rightarrow \infty} \frac{x_0}{c} = \frac{1}{v + \sqrt{Iw}} \quad \text{and} \quad \lim_{c \rightarrow 0} \frac{x_0}{c} = \frac{1}{v + Iw} \quad \text{for all } v, I, w. \quad (19)$$

We can see that F_A -fairness, in this case, allocates more of the fraction of capacity to sources of type 0 than proportional fairness, getting further away from proportional fairness as capacity increases, and exactly equalling it in the case of zero capacity.

We can also see that here F_A -fairness allocates less capacity than max-min fairness for any capacity. When capacity is large we can see from Equation (19) that the distribution to type 0 sources can be approximated by $\frac{c}{v+\sqrt{Iw}}$.

A graph of $\frac{x_0(c)}{c}$ for F_A -fairness alongside graphs for proportional and max-min fairness is shown in Figure 2 for the example when $\eta = 0.5, r_0 = 5, I = 2, v = 3, w = 2$). This graph is representative of any parameter settings.

5.3. F_B analysis

Lemma 5.2 *The F_B -fairness distribution for the parking lot scenario where $\vec{n} = (v, w, w, \dots)$ can be shown to be given by*

$$x_0 = \max \left(\frac{c}{v + Iw} - \frac{(I-1)wr_0}{\eta(v + Iw)}, 0 \right). \quad (20)$$

x_0 is strictly increasing in c . $\lim_{c \rightarrow \infty} x_0/c = \frac{1}{v+Iw}$. Thus, when $I = 1$, F_B -fairness' fraction of capacity is the same as that for proportional fairness (and max-min fairness). When $I > 1$, the fraction of capacity allocated is always less than proportional fairness.

In the limiting case, i.e. for very small capacity relative to the number of competing sources, F_B -fairness allocates zero to type 0 sources.

6. Verification by Simulation

In this section, we investigate the convergence of the average rate of the time series for the sources for small values of η and r_0 , and also for more TCP-like settings for the parameters. This is done both for the cases of rate proportional feedback and rate independent feedback.

We do this by simulation of the stochastic process in the parking lot scenario where $\vec{n} = (v, w, w, \dots)$. We don't verify the validity of the model, described in Section 2, in representing a real TCP in the case of rare negative feedback and equal round trip times.

6.1. Rate Proportional Feedback

We first verify that the convergence holds for small increments of η and r_0 . We then show that the series converges for TCP-like settings. More precisely, we show that in a regime of rare negative feedback, the average of the series converges to that expected from F_A -fairness for TCP-like settings of η and r_0 i.e. the distributed rates eventually oscillate around the value determined by F_A -fairness.

For the simulations, we use the family of $g_l(f, d, p)$ functions such that g_l is 0 when the link usage is less than dc , 1 when the link usage exceeds capacity available on the link and an increasing function from 0 to 1 given by $\left(\frac{f-d}{1-d}\right)^p$ for link usage between dc and c .

At the start of each simulation, each x_i is assigned a random number from a uniform distribution on $(0, c)$. At each iteration, the expectation, E_i for each source i is calculated. Then a random number is drawn from a uniform distribution on $(0, 1)$. If this number is greater than or equal to the calculated expectation, a value of $E_i = 0$ is assumed to have occurred, and x_i is linearly increased by r_0 . Otherwise, x_i is multiplicatively decreased by η . The system continues to evolve until the total average capacity allocated does not change by a given tolerance.

The available simulation parameters are η , r_0 , τ , I , v , w , d and p . For each chosen parameter set, the simulation is run four times, and the average of all four are calculated along with determined confidence intervals.

With linear increase/multiplicative decrease, the aggregate average rate allocated on a link will always be less than a link's nominal capacity c . Thus the sum of the average rates of all sources converges to a value, c' , below this nominal rate c . How close c' is to c is determined by the efficiency of the g_l function in maximising overall throughput.

So, for each source, we consider the proportion of its average rate that it has of c' . This value is what we refer to as the scaled average. We obtain the F_A fairness distribution from Equation (17).

Small Values of η and r_0

Here we consider values of $\eta = r_0 = 0.01$ and $\tau = 0.2$. We varied the parameters as follows: $I = 2, 5$, v and $w = 1, 2, 6, 12$, $c = 250, 625$, $d = 0, 0.5, 1$, and $p = 1, 2, 5, 10$. In all cases except when $d = 1$, we found the scaled average to converge to that expected from F_A -fairness, which can be seen in Figure 3, which includes error bars for 95% confidence.

When $d = 1$, the assumption of rare negative feedback no longer held because every source was receiving a large amount of negative feedback at the same time.

TCP-like parameter settings

Here we set $\eta = 0.5$, $\tau = 0.2$ and $r_0 = \frac{1}{\tau}$. We varied the parameters as in the previous case. As before, we found the scaled average to converge to that expected from F_A -fairness except for the case $d = 1$. This is illustrated by the scatter plot in Figure 4 for simulation values not including the $d = 1$ case. The error bars for 95% confidence are there, but perhaps not too visible given that the highest confidence interval is ± 0.002 .

To summarise, we have established that F_A -fairness is a realistic model for TCP-like connections with equal round-trip times.




Figure 3. F_A versus simulation results of type 0 sources' fraction of capacity. For small values of η and r_0 .




Figure 4. F_A versus simulation results of type 0 sources' fraction of capacity. For TCP-like parameter settings.




Figure 5. F_B versus simulation results of type 0 sources' fraction of capacity. For small values of η and r_0




Figure 6. F_B versus simulation results of type 0 sources' fraction of capacity. For TCP-like parameter settings.

6.2. Rate Independent Feedback Simulation

The case when the feedback is assumed to be rate independent as described in Section 2 was also simulated. This was done for small and TCP-like values of η and r_0 and the results compared with the values as determined by F_B -fairness.

We found that in both cases, the results agree with that anticipated from F_B -fairness, the main finding being that even with TCP-like parameter settings, the average rate converges in agreement with F_B .

We preserve the same conditions for simulation as in the rate proportional feedback case. The only difference is that the expectation of $E_i(t)$ is given by Equation (4) rather than Equation (3).

Small Values of η and r_0

Again we consider values of $\eta = r_0 = 0.01$, where $\tau = 0.2$ and for the same range of parameters as in the previous simulations. We found the scaled average to converge to that expected from F_B -fairness. This is shown in Figure 5.

TCP-like parameter settings

Here we set $\eta = 0.5$, $\tau = 0.2$ and $r_0 = \frac{1}{\tau}$. Again the same parameter set was used. Figure 6 shows the converged rate of x_0 sources versus results from calculating F_B -fairness.

Even when F_B -fairness determines that sources of type 0 should be allocated a rate of zero, the result converges to almost zero. This is in contrast to the rate allocated by proportional fairness. For a typical example, in one case the simulation average rate for type 0 sources converged about 0.0000006. Here, F_B would allocate 0 to type 0 sources, while proportional fairness would allocate 0.08333323.

7. Rate as a function of packet loss ratio

The analysis also provides a simple means to derive the source rates as a function of the packet loss ratio experienced by the source. For a given rate distribution vector \vec{x} , the packet loss ratio $q_i(t)$ over the path of source i is $q_i(t) = \sum_{l=1}^L g_l(f_l(\vec{x}(t)))A_{l,i}$ and we interpret Equation (3) by observing that, with case A, the expected feedback over one time cycle of duration τ is proportional to the number of packets sent $x_i(t)\tau$. With the hypothetical case B, we would say that the feedback is proportional to the packet loss ratio, but independent of the number of packets sent over one time interval (Equation (4)).

In the limit, we must have, for case A, $\lim_{t \rightarrow +\infty} \frac{dx_i(t)}{dt} = 0$ which, combined with Equation (9) gives $r_0 - \tau x_i^*(r_0 + \eta x_i^*)q_i^* = 0$ where $\lim_{t \rightarrow +\infty} x_i(t) = x_i^*$ and $\lim_{t \rightarrow +\infty} q_i(t) = q_i^*$.

Solving for x_i^* gives

$$x_i^* = \frac{-\tau q_i^* r_0 + \sqrt{4r_0 \tau q_i^* \eta + \tau^2 q_i^{*2} r_0^2}}{2\tau q_i^* \eta}. \quad (21)$$

For very small loss ratio q_i^* , the leading term in Equation (21) is given by $x_i^* \equiv_{q_i^* \rightarrow 0} \sqrt{\frac{r_0}{\tau q_i^* \eta}}$.

In the case of a TCP connection, we have $r_0 = \frac{1}{\tau}$ (packets per second) and $\eta = 0.5$. The previous equations give rates in packets per seconds; calling MSS the packet size in bits, we obtain the rates in bits per second from the previous equation:

$$x_i^* \equiv_{q_i^* \rightarrow 0} \frac{MSS}{\tau} \frac{C}{\sqrt{q_i^*}} \text{ b/s}$$

with $C = \sqrt{2}$. This last result is in line with a family of similar results [10,2,8]. Our results differs in the value of C , which we attribute to the fact that we have assumed a fluid model converging towards some equilibrium, whereas in reality the TCP window size oscillates around some equilibrium.

If we did the same analysis with the modelling of case B, we would find that the leading factor in x_i^* would be in $\frac{1}{q_i^*}$, which does not match the previous results. We interpret this as a further confirmation that model A is closer to reality than model B.

8. Conclusions and Future Work

TCP compliant sources with equal round trip times competing for bandwidth do not, as was previously thought, end up with a distribution of rates in accordance with proportional fairness.

Rather, we show that when feedback is rate dependent and negative feedback rare, the distribution agrees with F_A -fairness. In addition, we confirm this by derivation of the standard TCP throughput as a function of loss formula.

Even in the cases where feedback could no longer be assumed to be rate dependent, we have shown that proportional fairness would only approximate the long term rate distribution, and would be reflected closer by F_B -fairness.

An assumption of rare negative feedback is valid when the increments are small (i.e. the round-trip time τ is small) and the losses relatively low. It is our belief that these results essentially hold when we remove the assumption of rare negative feedback, but this remains to be verified. The larger puzzle will be solved when the rate distribution behaviour is determined for different round trip times and this forms part of our intended ongoing work.

It is known that TCP gives less throughput to connections with longer round trip times. Based on our analysis there are two possible reasons: F_A -fairness which provides less to connections that use several hops; or the fact that TCP maintains a sending window rather than a sending rate. It is not clear to us what the respective affects of each of these factors are.

The results shown have potential implications for multimedia applications which are and will be expected (or even required) to be "TCP friendly" conformant [12]. Namely, they behave like TCP source would in receipt of both negative and positive feedback.

REFERENCES

1. A. Benveniste, M. Metivier, and P. Priouret. *Adaptive Algorithms and Stochastic Approximations*. Springer Verlag, Berlin, 1990.
2. S. Floyd. Connections with multiple congested gateways in packet switched networks, part 1. *ACM Computer Communication Review*, 22(5):30–47, Oct 1991.
3. S. Floyd, V. Jacobson, Random Early Detection gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking*, V.1 N.4, August 1993, p.397-413.
4. S. Golestani and S. Bhattacharyya. End-to-end congestion control for the internet: A global optimization framework. *Proc of ICNP, Oct 98*, 1998.
5. F.P. Kelly, A. K. Maulloo, D.K.H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49, 1998.
6. H. J. Kushner and D.S. Clark. Stochastic approximations for constrained and unconstrained systems. *Applied Mathematical Sciences*, 26, 1978.
7. Liung L. Analysis of recursive stochastic algorithms. *IEEE Trans. on Automatic Control*, 22:551–575, 1977.
8. T. V. Lakshman and U. Madhow. The performance of TCP for networks with high bandwidth delay products and random loss. *IEEE/ACM Trans on Networking*, 5(3):336–350, June 1997.
9. L. Massoulie and J. Roberts. Fairness and QoS for elastic traffic. *CNET*, 1998.
10. M. Mathis, J. Semke, J. Mahdavi, T. Ott. The macroscopic behaviour of the TCP congestion avoidance algorithm. *Comp, Comm. Review*, 3, July 1997.
11. R.K. Miller, A. N. Michell. *Ordinary Differential Equations*. Academic Press, 1982.
12. TCP friendly web site. http://www.psc.edu/networking/tcp_friendly.html
13. D. Chiu, R. Jain Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks. *Computer Networks and ISDN Systems*, vol. 17, pp. 1-14, June 89.